

ABSTRACT

SG'7

Roughly described, the interpretation phase of a system using an Adaptive Agent-Oriented Software Architecture allows queried agents to respond to queries before they have all their own responses from their own downchain agents. In one embodiment, queried agents respond at a fixed time after receipt of a query, whether or not they have received all responses from their own downchain agents. In another embodiment, a queried agent makes claims to its upchain inquiring agent in response to each claim that the queried agent receives from its own downchain agents. In another embodiment, a queried agent can receive a particular query more than once, and in response to each, the agent responds with whatever claims it then has. In order to limit the duration of time during which queries are active in the network, and thus during which new claims can still be made, the agent originating a query can send a "forget-problem" or a "commit" message down into the network after some period of time. Alternatively or additionally, the originating agent can include a "depth-of-search" indication with each query, thereby preventing propagation of the query through more than the indicated number of agents. In the latter alternative, the originating agent can subsequently make the same query to the same downchain agents, but with an increased depth-of-search indication, if for example the originating agent is not yet satisfied with the claims it received in response to the first query.